

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.

THIS PAGE BLANK (USPTO)

Single-instruction multiple-data is a new class of integrated video signal processors especially suited for real-time processing of two-dimensional images. The single-instruction, multiple-data architecture is adopted to exploit the high degree of parallelism inherent in many video signal processing algorithms. Features have been added to the architecture which support conditional execution and sequencing - an inherent limitation of traditional single-instruction multiple-data machines. A separate transfer engine offloads transaction processing from the execution core, allowing balancing of input/output and compute resources - a critical factor in optimizing performance for video processing. These features, coupled with a scalable architecture allow a united programming model and application driven performance.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

| | | | | | |
|----|--------------------------|----|------------------------------------------|----|--------------------------|
| AT | Austria | FR | France | MR | Mauritania |
| AU | Australia | GA | Gabon | MW | Malawi |
| BB | Barbados | GB | United Kingdom | NL | Netherlands |
| BE | Belgium | GN | Guinea | NO | Norway |
| BF | Burkina Faso | GR | Greece | NZ | New Zealand |
| BG | Bulgaria | HU | Hungary | PL | Poland |
| BJ | Benin | IE | Ireland | PT | Portugal |
| BR | Brazil | IT | Italy | RO | Romania |
| CA | Canada | JP | Japan | RU | Russian Federation |
| CF | Central African Republic | KP | Democratic People's Republic of Korea | SD | Sudan |
| CG | Congo | KR | Republic of Korea | SE | Sweden |
| CH | Switzerland | LI | Liechtenstein | SK | Slovak Republic |
| CI | Côte d'Ivoire | LK | Sri Lanka | SN | Senegal |
| CM | Cameroon | LU | Luxembourg | SU | Soviet Union |
| CS | Czechoslovakia | MC | Monaco | TD | Chad |
| CZ | Czech Republic | MG | Madagascar | TG | Togo |
| DE | Germany | ML | Mali | UA | Ukraine |
| DK | Denmark | MN | Mongolia | US | United States of America |
| ES | Spain | | | VN | Viet Nam |
| FI | Finland | | | | |

DATA PROCESSING SYSTEM

Background Of The Invention1) Field of the Invention.

This invention relates to the field of video signal processing, and, in particular, to video signal processing using an architecture having a plurality of parallel execution units.

2) Background Art

It is well known in the prior art to use multiple-instruction multiple-data systems for video signal processing. In a multiple-instruction multiple-data execution of an algorithm each processor of the video signal processor may be assigned a different block of image data to transform. Because each processor of a multiple-instruction multiple-data system executes its own instruction stream, it is often difficult to determine when individual processors have completed their assigned tasks. Therefore, a software

synchronization barrier may be used to prevent any processors from proceeding until all processors in the system reach the same point. However it is sometimes difficult to determine where synchronization barriers are required. If a necessary barrier is omitted by a user then the resulting code may be nondeterministic and re-execution of the code on the same data may yield different results.

An alternate architecture known in the prior art is single-instruction multiple-data architecture. Single-instruction, multiple-data is a restricted style of parallel processing lying somewhere between traditional sequential execution and multiple-instruction multiple-data architecture having interconnected collections of independent processors. In the single-instruction, multiple-data model each of the processing elements, or datapaths, of an array of processing elements or datapaths executes the same instruction in lock-step synchronism. Parallelism is obtained by having each datapath perform the same operation on a different set of data. In contrast to the multiple-instruction, multiple-data architecture, only one program must be developed and executed.

Referring now to Fig. 1, there is shown prior art single-instruction multiple-data architecture 100.

A conventional single-instruction multiple-data system, such as architecture 100, comprises a controller 112, a global memory 126 and execution datapaths 118a-n. A respective local memory 120a-n may be provided within each execution datapath 118a-n. Single-instruction multiple-data architecture 100 performs as a family of video signal processors 118a-n united by a single programming model.

Single-instruction multiple-data architecture 100 may be scaled to an arbitrary number n of execution datapaths 118a-n provided that all execution datapaths 118a-n synchronously execute the same instructions in parallel. In the optimum case, the throughput of single-instruction multiple-data architecture 100 may theoretically be n times the throughput of a uniprocessor when the n execution datapaths 118a-n operate synchronously with each other. Thus, in the optimum case, the execution time of an application may be reduced in direct proportion to the number n of execution datapaths 118a-n provided within single-instruction multiple-data architecture 100. However, because of overhead in the use of execution datapaths 118a-n, this optimum is never reached.

Architecture such as single-instruction multiple-data architecture 100 works best when

executing an algorithm which repeats the same sequence of operations on several independent sets of highly parallel data. For example, for a typical image transform in the field of video image processing, there are no data dependencies among the various block transforms. Each block transform may be computed independently of the others.

Thus the same sequence of instructions from instruction memory 124 may be executed in each execution datapath 118a-n. These same instructions are applied to all execution datapaths 118a-n by way of instruction broadcast line 116 and execution may be independent of the data processed in each execution datapath 118a-n. However, this is true only when there are no data-dependent branches in the sequence of instructions. When data-dependent branches occur, the data tested by the branch will, in general, have different values in each datapath. It will therefore be necessary for some datapaths 118a-n to execute the subsequent instruction and other datapaths 118a-n to not execute the subsequent instruction. For example, the program fragment of Table I clips a value v between a lower limit and an upper limit:

```
local v;  
  
...  
  
v = expression  
if (v > UPPER_LIMIT)  
    v = UPPER_LIMIT;  
if (v < LOWER_LIMIT)  
    v = LOWER_LIMIT;
```

Table I

The value being clipped, v, is local to each execution datapath 118a-n. Thus, in general, each execution datapath 118a-n of single-instruction multiple-data architecture 100 executing the program fragment of Table I may have a different value for v. In some execution datapaths 118a-n the value of v may exceed the upper limit, and in others v may be below the lower limit. Other execution datapaths 118a-n may have values that are within range. However the execution model of single-instruction multiple-data architecture 100 requires that a single identical instruction sequence be executed in all execution datapaths 118a-n.

Thus some execution datapaths 118a-n may be required to idle while other execution datapaths 118a-n perform the conditional sequence of Table I. Furthermore, even if no execution datapaths 118a-n of single-instruction multiple-data architecture 100 are required to execute the conditional sequence of the

program fragment of Table I, all execution datapaths 118a-n would be required to idle during the time of the conditional sequence. This results in further inefficiency in the use of execution datapaths 118a-n within architecture 100.

Another problem with systems such as prior art single-instruction multiple-data architecture 100 is in the area of input/output processing. Even in conventional uniprocessor architecture a single block read instruction may take a long period of time to process because memory blocks may comprise a large amount of data in video image processing applications. However, this problem is compounded when there is a block transfer for each enabled execution datapath 118a-n of architecture 100 and datapaths 118a-n must compete for access to global memory 126. For example, arbitration overhead may be very time consuming.

The alternative of providing each execution datapath 118a-n with independent access to external memory 126 is impractical for semiconductor implementation. Furthermore, this alternative restricts the programming model so that data is not shared between datapaths 118a-n. Thus further inefficiency results due to the suspension of processing of instructions until all the block reads are

completed. This may be seen in the discrete cosine transform image kernel of Table II:

```
for (i = 0; i < NUMBEROFBLOCKS; i = i + 4) {  
    k = i + THIS_DP_NUMBER;  
    read_block(original_image[k], temp_block);  
    DCT_block(temp_block);  
    write_block(xform_image[k], temp_block);  
};
```

Table II

The read_block and write_block routines of the instruction sequence of Table II must be suspensive. Each must be completed before the next operation in the kernel is performed. For example, read_block fills temp_block in local memory 120a-n with all of its local values. These local values are then used by DCT_block to perform a discrete cosine transform upon the data in temp_block. Execution of the discrete cosine transform must wait for all of the reads of the read_block command of all execution datapaths 118a-n to be completed. Only then can the DCT_block and write_block occur. Thus, by the ordering rules above, read_block must be completed before the write_block is processed, or the DCT_block is executed.

Referring now to Fig. 2, there is shown processing/memory time line 200. The requirements imposed by the ordering rules within single-instruction

multiple data architecture 100 result in the sequentialization of memory transactions and processing as schematically illustrated by processing/memory time line 200. In time line 200, memory read_block time segment 202 of execution datapath 118a-n must be completed before processing of DCT_block time segment 204 may begin. Processing DCT_block time segment 204 must be completed before memory write_block time segment 206 may begin. Only when memory write_block time segment 206 of a execution datapath 118a-n is complete, can memory read_block time segment 208 of a execution datapath 118a-n begin. Execution and access by second execution datapath 118a-n is sequentialized as described for the first.

This problem occurs in high performance disk input/output as well. In a typical disk input/output operation an application may require a transfer from disk while continuing to process. When the data from disk are actually needed, the application may synchronize on the completion of the transfer. Often, such an application is designed to be a multibuffered program. In this type of multibuffered program, data from one buffer is processed while the other buffer is being filled or emptied by a concurrent disk transfer. In a well designed system the input/output time is com-

pletely hidden. If not, the execution core of single-instruction multiple-data architecture 100 is wait-stated until the data becomes available. This causes further degrading of the performance of the single-instruction multiple-data architecture 100.

Summary of the Invention

A single-instruction, multiple-data image processing system is provided for more efficiently using parallel datapaths when executing an instruction sequence having conditionals. At least two conditionals are sequentially determined according to the instructions. A respective mask flag is set for each conditional, wherein the mask flag is effective to instruct the datapath whether to execute an instruction or to idle during a selected instruction cycle. The mask flags are sequentially stored and later retrieved in a predetermined order. The execution unit of the datapath determines whether to execute an instruction or idle during selected instruction cycles according to the mask flags which are sequentially retrieved. Each datapath of the image processing system has an execution unit and a local memory which may be accessed only by the execution unit in the same datapath as the local memory. Access between the execution unit and

the local memory is by way of one port of a dual-ported memory.

Thus the system of the present invention solves several problems associated with the single-instruction, multiple data architecture. One problem solved by the architecture of the present invention is that of efficiently permitting some datapaths, but not others, to execute particular instructions which have been issued by a common instruction store. This problem arises as a result of data dependencies and is solved by conditional execution and execution masks. During conditional execution every instruction which has been issued is executed or not executed by a particular datapath depending on the state of a data dependent condition code as calculated by the particular datapath. Each datapath can also set or clear its own particular local execution mask depending on the state of the data dependent condition code as calculated by the particular datapath. If the execution mask for a particular data path is active, the datapath ignores any instruction which has been issued. A special instruction command is provided to reactivate idle datapaths.

The execution mask feature is more general than the conditional execution feature because it

permits nesting. The execution masks are saved on a stack within the datapaths. If the active subset of datapaths encounter data dependencies, additional datapaths may be turned off. The earlier state of the processor is restored by popping the execution mask stack. The conditional execution feature complements the execution mask feature by permitting a very efficient treatment of simple cases in which only a few instructions are data dependent and no nesting is involved.

All transfer between the local memory and the global memory take place using one dedicated port of the local memory. The transfers are scheduled and controlled by a common unit called the block transfer controller. The block transfer controller, along with the dedicated port of the dual ported local memory, permit each access to global memory by a datapath to be overlapped with its instruction processing. This usually avoids stalling the processor.

Brief Description of the Drawings

Fig. 1 shows a block diagram representation of a prior art single-instruction multiple-data architecture suitable for processing highly parallel data such as data representative of video images;

Figs. 2 shows a processing/memory time line for the image processing architecture of Fig. 1;

Fig. 3 shows the single-instruction multiple-data architecture image processor of the present invention;

Figs. 4A-D show execution mask stacks;

Fig. 5 shows a processing/memory time line for the architecture of Fig. 3;

Fig. 6 shows a schematic representation of a memory-to-memory block transfer within the architecture of Fig. 3;

Fig. 7 shows a linked request list formed of command templates linked to each other and containing the parameters required to specify a memory-to-memory block transfer such as the transfer represented by Fig. 6;

Fig. 8 shows a simplified alternate embodiment of the single-instruction multiple-data architecture image processor of Fig. 3, wherein there is provided a four-execution datapath architecture;

Fig. 9 shows a block diagram representation of a system for performing internal scalar transfers

between the datapaths of the single-instruction multiple-data image processor of Fig. 3;

Fig. 10 shows a block diagram representation of a statistical decoder for decoding variable length codes within the single-instruction multiple-data image processor of Fig 3;

Fig. 11 shows a binary decoding tree for decoding variable length codes within the statistical decoder of Fig 10.

Detailed Description Of The Invention

Referring now to Fig. 3, there is shown single-instruction multiple-data architecture image processor 300 of the present invention. Single-instruction multiple-data image processor 300 is provided with execution masks and conditional control flow during conditional branches in order to provide more efficient use of computation time in execution datapath 358a-n of image processor 300. Each of these two mechanisms addresses one of two distinct control needs within image processor 300.

Each execution datapath 358a-n of a single-instruction multiple-data image processor 300 is provided with a respective execution unit 360a-n and

local memory 362a-n. Each execution unit 360a-n of execution datapath 358a-n is coupled to its respective local memory 362a-n by way of a respective local memory port 361a-n and to system memory 364 and global memory 366 by way of a respective global memory port 363a-n. Local memory ports 361a-n and global memory ports 363a-n, together, provide each execution datapath 358a-n with a dual port architecture to permit each execution unit 360a-n to access its respective local memory 362a-n simultaneously with data transfer between local memories 362a-n and memories 364, 366. It will be understood that within the dual port architecture of image processor 300, no execution unit 360a-n may directly access any local memory 362a-n except its own.

During execution of instructions, instruction sequence controller 352 of single-instruction multiple-data image processor 300 simultaneously applies the same instruction to every execution datapath 358a-n by way of broadcast instruction line 356. The instructions applied by sequence controller 352 are previously stored in either system memory 364 or global memory 366. The instructions received by sequence controller 352 are applied to sequence controller 352 by way of memory instruction line 356. Within image processor 300, conditional execution permits each

datapath 358a-n to execute or not execute a particular issued instruction depending on the state of the local datapath condition flag. Hardware execution masks, residing within execution units 360a-n of image processor 300, permit individual datapaths 358a-n to turn off execution of a sequence of issued instructions for an arbitrary period of time. These two mechanisms decrease the amount of wait stating or idling of execution datapaths 360a-n within single-instruction multiple-data image processor 300, thereby permitting more efficient use of execution datapaths 358a-n.

Control over whether an instruction issued by sequence controller 352 is executed or ignored by an individual execution datapath 358a-n is required for data-dependent computation in a single-instruction multiple-data architecture such as the architecture of image processor 300. It is required because each execution datapath 358a-n may have a different value when a test is performed as part of a conditional branch. Thus each execution datapath 358a-n within image processor 300 of the present invention is provided with individual datapath execution masks.

It is equally important to control the sequence of instructions provided by sequence controller 352 to execution datapaths 358a-n by way of

broadcast instruction line 356. This control is essential for loops and may also be used to optimize data-dependent execution wherein no execution datapath 358a-n is required to execute a conditional sequence of instructions.

For the purpose of executing a conditional branch within image processing architecture 300, each datapath 358a-n tests the condition of a conditional branch and independently sets its own flags according to its own local determination. Signals representative of these flags are applied by each execution datapath 358a-n to instruction sequence controller 352 by way of flag signal lines 354a-n.

Rather than automatically wait-stating all execution datapaths 358a-n during a conditional, single-instruction multiple-data architecture 300 of the present invention uses the flag signals of flag lines 354 to apply a consensus rule. In the consensus rule of image processor 300, sequence controller 352 does not apply a conditionally executed instruction sequence to broadcast instruction line 356 unless flag lines 354 signal controller 352 that every execution datapath 358a-n requires the instruction sequence. This prevents the inefficiency which results when some execution datapaths 358a-n are wait-stated for the

duration of a sequence which is not executed by some of the datapaths 358a-n.

Both mechanisms, conditional execution and execution masks may be used to implement the conditional execution within image processor 300, when some but not all datapaths 358a-n require it. Of these two mechanisms, execution masks EM are more the general. The execution mask flag is appended to the normal set of local arithmetic condition code flags within each execution unit 360a-n. When an execution mask flag EM is set within an execution unit 360a-n of execution datapaths 358a-n and sequence controller 352 applies the conditional sequence to broadcast instruction line 356, each execution unit 360a-n having its execution mask flag EM set ignores the instructions.

The only exceptions to instructions being ignored by execution datapath 358a-n within image processor 300 when an execution mask flag EM is set are 1) the instruction which restores the state of the previous execution mask flag, and 2) those instructions which unconditionally modify the execution mask flag EM. These instructions are executed by all execution unit 358a-n even if the execution mask flag EM within a datapath 358a-n is set. Thus, if the execution mask

flag EM is set in a selected execution unit 360a-n, instructions from instruction sequence controller 352 are ignored by the selected execution unit 360a-n.

It is then possible to encode a conditional thresholding program fragment within single-instruction multiple-data architecture image processor 300 using execution masks EM. This thresholding program is set forth in the instruction sequence of Table III. The instruction sequence of Table III is adapted to perform, within image processor 300, the clipping operation performed by the instruction sequence of Table I within prior art architecture 100. In this instruction sequence a local value v is constrained within a range between the values LOWER_LIMIT and UPPER_LIMIT.

| | |
|--------------------|------------------------------------------------------------------------------------------------------|
| CMP v,UPPER_LIMIT; | compare and set flags on all execution datapaths 358a-n |
| MOV EM,LE; | set execution masks EM on execution datapaths 358a-n with less than or equal flag set |
| MOV v,UPPER_LIMIT; | update v only on execution datapaths 358a-n with greater than flags set |
| MOV EM,0; | every execution datapath 358a-n executes unconditional reset of EMs, activating all datapaths 358a-n |
| CMP v,LOWER_LIMIT; | compare and set flags on all execution datapaths 358a-n |
| MOV EM,GE; | set EM on datapaths 358a-n with greater than or equal flag set |
| MOV v,LOWER_LIMIT; | update v on execution datapaths 358a-n with less than flag set |
| MOV EM,0; | reenable every execution datapath 358a-n |

Table III

The first instruction of Table III, executed by all execution datapaths 358a-n of single-instruction multiple-data architecture image processor 300, compares the local value of v for each execution datapath 358a-n against the same upper threshold UPPER_LIMIT. The MOV EM,LE instruction of Table III is then executed by all execution datapaths 358a-n of image processor 300. A respective execution mask flag EM is thereby determined within each execution datapath 358a-n according to the comparison of the local v. Each execution datapath 358a-n is thus provided with a

respective setting of the flag EM in its individual flag register.

In execution datapaths 358a-n where the less-than-or-equal condition is met, the MOV EM,LE instruction results in the execution mask flag EM being set to the value one. Execution datapaths 358a-n wherein the execution mask flag EM is set to the value one are disabled. These disabled execution datapaths 358a-n ignore instructions applied by sequence controller 352 by way of broadcast instruction line 356. In particular, execution datapaths 358a-n having their execution mask flag EM set to one by the MOV EM,LE instruction of Table III ignore the MOV v,UPPER_LIMIT instruction. In execution datapaths 358a-n where the less-than-or-equal condition is not met, the execution mask flag is set to the value zero. These datapaths 358a-n execute the MOV v, UPPER_LIMIT instruction, thereby clipping any local values of v which were greater than UPPER_LIMIT.

Thus the MOV v,UPPER_LIMIT, is executed only by those execution datapaths 358a-n where the greater-than condition was met in the first instruction of Table III and the execution mask flag EM has the value zero. The fourth instruction of Table III, MOV EM,0, unconditionally resets the execution mask flag EM of

all execution datapaths 358a-n of image processor 300 including those execution datapaths 358a-n wherein the execution mask flag EM was set by the CMP v,UPPER_LIMIT instruction. Thus the update of the value of v within the sequence of Table III occurs only in those execution datapaths 360a-n in which it is required and execution mask flags EM of all execution units 360a-n are reset to zero for the next compare instruction.

The example of Table III may be simplified using conditional execution as illustrated by the instruction sequence of Table IV.

| | |
|----------------------------|------------------------------------------------------------------------------------------------|
| CMP v,UPPER_LIMIT | Compare local values of v and set all flags in all execution datapaths 358a-n |
| IF (GT): MOV v,UPPER_LIMIT | Executed only by execution datapaths 358a-n where local values of v exceed UPPER_LIMIT. |
| CMP v,LOWER_LIMIT | Set all flags on all execution datapaths 358a-n and compare. |
| IF (LT): MOV v,LOWER_LIMIT | Executed only by execution datapaths 358a-n where local values of v are less than LOWER_LIMIT. |

Table IV

Conditional execution within image processor 300 permits every instruction to be executed by each datapath 358a-n based on the data-dependent condition code local to the individual datapath 358a-n. The conditional execution feature of image processor 300 is thus more efficient than prior art architecture, such as prior art architecture 100. However, the method described does not allow for nesting of data dependent execution when using execution masks EM. In these nested execution cases conditional execution may still be used for improved efficiency in the innermost of the nested data dependencies. The two mechanisms are therefore complementary and may be used together to achieve maximum efficiency.

In order to permit nesting, the execution mask is generalized to a stack. In the preferred embodiment of image processing system 300, the execution mask stacks are respective hardware stacks residing in each execution unit 358a-n. The push command pushes the local execution mask or condition code of an individual execution datapath 358a-n onto its individual execution mask stack. Push and pop operations are executed by all execution datapaths 358a-n, regardless of whether they are active.

In the case of an inactive datapath 358a-n, the condition code pushed onto the stack has the value one, indicating that the inactive datapath 358a-n is off. The remaining active datapaths 358a-n execute the compare against UPPER_LIMIT. The subset of datapaths 358a-n not requiring a clip against UPPER_LIMIT are turned off by the next push operation. Following the clip, a pop restores the prior state and those datapaths 358a-n with clipping enabled are all reenabled for the test against LOWER_LIMIT. Following the similar clipping against LOWER_LIMIT, the final pop operation reenables all datapaths 358a-n. It is necessary for push and pop to execute in all datapaths 358a-n in order to insure the consistency of the execution mask stack.

In the case of a nested conditional, a conditional sequence expression is executed only when some further condition is true. For example, in some applications a determination whether v is within range, similar to the determination of Table V, may be made only if the clipping routine is enabled. This is indicated when the variable enable_clipping is non-zero in the instruction sequence of Table V.

```

local v;
...
v = expression)
if (enable_clipping){
    if (v > UPPER_LIMIT)
        v = UPPER_LIMIT;
    if (v < LOWER_LIMIT)
        v = LOWER_LIMIT;
}
...

```

Table V

However, when executing the instruction sequence of Table V it is not possible to merely compare enable_clipping in each execution datapath 358a-n and set the execution mask flag EM accordingly when a MOV EM,0 instruction corresponding to the upper limit test is executed. Because such a setting of the execution mask flag EM would be unconditional, all execution datapaths 358a-n would execute it. This would cause all execution datapaths 358a-n within image processor 300 to become enabled.

Thus, even execution datapaths 358a-n, where enable_clipping was false, would be enabled. Therefore, all execution datapaths 358a-n would perform the subsequent lower limit test, even those that should not perform any clipping of v at all because their clipping routine was not enabled. However, conditional execution and execution masks can both be used to efficiently implement an enabled clipping operation.

Additionally, it may be implemented without the use of conditional execution. This is useful for illustrating the generality of the execution mask technique which may be applied to arbitrary levels of nesting.

Conditional execution and execution masks may both be used to efficiently implement the example with enabled clipping as shown in Table VI.

| | |
|----------------------------|-----------------------------------------------------------------------------------------------|
| CMP ENABLE,0 | |
| MOV EM,EQ | Disable execution datapaths 358a-n where clipping is not enabled |
| CMP v,UPPER_LIMIT | Compare and set all flags on all execution datapaths 358a-n |
| IF (GT): MOV v,UPPER_LIMIT | Executed only by execution datapaths 358a-n where local values of v exceed UPPER_LIMIT |
| CMP v,LOWER_LIMIT | Compare and set all flags on all execution datapaths 358a-n |
| IF (LT): MOV v,LOWER_LIMIT | Executed only by execution datapaths 358a-n where local values of v are less than LOWER_LIMIT |
| MOV EM,0 | Reenable all execution datapaths 358a-n |

Table VI

This example can also be implemented without the use of conditional execution as shown in Table VII.

This illustrates the generality of the execution mask technique which may be applied to arbitrary levels of nesting.

```

CMP ENABLE,0

PUSH EM,EQ

    CMP v,UPPER_LIMIT
    PUSH EM,LE
        MOV v,UPPER_LIMIT
    POP

    CMP v,LOWER_LIMIT
    PUSH EM,GE
        MOVE v,LOWER_LIMIT
    POP

POP

```

Table VII

Referring now to Figs. 4A-D, there are shown execution mask stack maps 402, 404, 406, and 408, representing execution mask stacks 359a-n. Stack maps 402, 404, 406, and 408 schematically illustrate portions of execution units 360a-n within execution datapath 358a-n of image processor 300. As previously described, in the preferred embodiment of execution units 360a-n, dedicated hardware is provided within execution units 360a-n to perform the functions of stacks 359a-n. Execution mask flag stack 359i of datapath 358x, where $1 \leq x \leq n$, contains y execution mask flags, where execution mask flag EMx.1 is the

first execution mask pushed onto mask stack 359x, EMx.2 is the second mask pushed onto execution mask flag stack 359x, and so on.

In general, an expression that sets an execution mask flag EMx.y may be coded to work correctly in any nested conditional, such as the conditional in the instruction sequence of TABLE VI, by saving the state of execution mask flag EMx.y in a temporary location at the beginning of the expression and restoring it at the end. Execution mask flag stacks 359a-n are provided for this purpose within respective execution units 360a-n of datapaths 358a-n.

It will be understood by those skilled in the art that the index x identifies a particular execution datapath 358a-n having an execution mask stack 359x within execution unit 360x and that the index y represents the number of masks stacked in execution mask stack 359x. Thus single-instruction multiple-data image processor 300 of the present invention provides execution mask flag stacks 359a-n for sequentially storing and sequentially retrieving execution mask flags EMx.y within each execution unit 360a-n of each execution datapath 358a-n.

When the first instruction of the sequence of Table VII, CMP ENABLE,0, is executed, the equal flag of execution datapath 358a-n and the equal flag of execution datapath 358n are set because the user has defined them to be enabled. Execution masks EMa.1, EMb.1, and EMn.1, appear in mask stacks 359a,b,n respectively within execution units 360a,b,n when the next instruction, PUSH EM,EQ, is executed by datapaths 358a,b,n.

Because datapath 358a and datapath 358n of execution mask stack map 402 are enabled, execution mask EMa.1 in execution mask stack 359a and execution mask EMn.1 in execution mask flag stack 359n have a value of zero. This permits execution datapaths 358a,n to execute instructions from sequence controller 352. Execution mask EMb.1, the first mask in execution mask stack 359b, has the value one because the user has previously defined datapath 358b to be disabled. Because execution datapath 358b is disabled, it does not execute instructions from sequence controller 352.

The next instruction of the instruction sequence of Table VII, CMP v,UPPER_LIMIT, is a nested data dependency. It causes some datapaths 358a-n which execute it to set the LE flag according to the local value of v. Other datapaths 358a-n which execute this

instruction do not set the LE flag. Furthermore, some datapaths 358a-n do not execute the instruction at all. For example, execution datapath 358b does not execute the instruction because the top most mask stacked in execution mask stack 359b, EMB.1, has a value of one. Execution datapath 358a and execution datapath 358n do execute the comparison instruction because their top most execution masks, EMA.1 and EMn.1, both have the value of zero.

When execution datapaths 358a,n execute the comparison and the instruction PUSH EM,LE, they each push a new execution mask, EMA.2 and EMn.2. These new execution masks are stored in their respective execution mask stacks 359a,n of the new execution mask stack map 404. A further execution mask EMB.2, having the value one, appears on stack 359b within execution unit 360b of datapath 358b. Because datapath 353b is inactive the disabling execution mask is merely reproduced by the push operation.

For the purposes of illustration, consider the case wherein the results of the upper limit comparison are such that execution datapath 358a must clip its local value of v while the local value of v of execution datapath 358n is not above the range and does not require clipping. In this case execution mask flag

EMa.2, pushed onto mask stack 359a, has the value of zero, and execution mask EMn.2, pushed onto mask stack 359n has the value one. Thus, during the execution of the next instruction, wherein the upper limit is moved to the local value of v, execution datapath 338a is active and performs the move, thereby clipping the value of v local to datapath 358a. During this instruction cycle execution datapaths b,n are inactive and do not perform the move, although they were inactivated at different points within the instruction sequence of TABLE VII.

During the execution of the next instruction, POP, execution masks EMa.2, EMb.2, and EMn.2 are removed from mask stacks 359a,b,n respectively as shown in execution mask stack map 406. The top most execution mask in mask stack 359n, EMn.1, has a value of one as previously determined by the enable comparison. Thus, by stacking execution masks EMn.1 and EMn.2 within stack 359n of execution datapath 358n, execution of datapath 358n was disabled during a nested loop and reenabled at the end of the nested loop. More generally, all execution datapaths 358a-n within single instruction multiple data video processor architecture 300 are able to idle or inactivate according to local data dependencies during a nested loop. Upon leaving

the nested loop, each execution datapath 358a-n may restore its execution mask status to its status prior to entering the loop. The POP instruction which follows then clears all execution masks EMI.1, as pushed onto execution mask stacks 359a-n by the PUSH instruction as shown in execution mask stack map 408.

It will be understood that the example of Table VII, as illustrated by execution mask stack maps 402, 404, 406, and 408, may occur embedded within a further instruction sequence (not shown). Instructions within the further instruction sequence may have pushed a plurality of flags onto mask stacks 359a-n previous to the first PUSH instruction of Table VII. Therefore the final POP instruction of Table VII may restore a previous execution mask status for execution datapaths 358a-n for further execution by image processor 300.

Each execution mask flag stack 359a-n within its respective execution datapath 358a-n thus provides automatic storage for execution mask flags EMx.y in image processor 300 of the present invention. In this method, execution mask flags EMx.y are pushed onto execution mask flag stacks 359a-n of execution datapaths 358a-n. When execution mask flags EMx.y are pushed onto stacks 359a-n, an operation may be performed by selected execution units 360a-n, and

execution mask flags $EMx.y$ may then be popped off stacks 359a-n. Stacks 359a-n containing execution masks $EMx.y$ within each execution datapath 358a-n of single-instruction multiple-data architecture 300 may have any number of entries. For example, execution datapaths 358a-n of image processor 300 may be provided with stacks 359a-n having sixteen or thirty-two entries.

The execution mask discipline described thus far provides a way to control, within an individual execution datapath 358a-n, the conditional execution of an instruction issued by sequence controller 352 of single-instruction multiple-data architecture 300. However, this execution mask discipline does not provide a way to conditionally control the sequence of instructions issued by sequence controller 352 during a conditional branch. For example:

```
local j;  
for (j=0 j<NUMBEROFBLOCKS; j=j+4){  
...  
}
```


The following instruction sequence of Table VIII performs this operation.

| | | |
|-----|-----------------------|-------------------------------|
| | MOV j,0; | initialize induction variable |
| L1: | CMP j,NUMBEROFBLOCKS; | test for end condition |
| | JGE L2 | exit if condition met by all |
| | | active datapaths 358a-n |
| | MOV EM,GE | turns off those datapaths |
| | | 358a-n |
| | | meeting exit condition |
| | | |
| | ADD j,4; | increment j |
| | JMP L1; | go back for more |
| L2. | | finished |

Table VIII

In the instruction sequence of Table VIII a copy of the local loop induction variable *j* exists in all execution datapaths 358a-n. The operation CMP *j*,NUMBEROFBLOCKS individually sets the execution mask flags EM_{x.y} of each execution unit 360a-n according to the local value of *j*. Because all execution datapaths 358a-n initialize *j* to the same value of zero and perform the same operation, ADD *j*,4, upon *j*, the execution mask flags EM of all execution datapaths 358a-n should be identical.

Since each execution datapath 358a-n may have a different number of blocks to process, the value NUMBEROFBLOCKS may vary from one execution datapath 358a-n to another. The instruction sequence for the loop is executed only if at least one datapath has an

index j less than NUMBEROFBLOCKS. Prior to executing the loop, those execution datapaths 358a-n which meet the exit condition are turned off by MOV EM,GE. When only a single execution datapath 358a-n is enabled within image processor 300, the enabled execution datapath 358a-n behaves like a conventional uniprocessor.

The consensus rule of image processor 300 allows the easy coding of a conditional program fragment that may be jumped over by all execution datapaths 358a-n if no execution datapath 358a-n requires the execution of the fragment. For example, occasionally v is negative. A very complex calculation requiring a great deal of processing time is required within execution units 360a-n when v is negative. If the calculation required for a negative v is sequenced, even if no execution datapath 358a-n requires it, extremely inefficient use is made of the processing power within single-instruction multiple-data architecture 300. Therefore, using the consensus rule, sequence controller 352 does not apply the instructions required for a negative v to broadcast instruction line 356 for transmission to execution datapaths 358a-n if no datapath 358a-n is processing a v having a negative value.

It is important to note that the consensus rule is complete. The dual to branching if all execution datapaths 358a-n satisfy the condition code ("if all") is to branch if any execution datapath 358a-n satisfies the condition code ("if any"). It is illustrated below how the branch "if any" function may be implemented using the branch "if all" type of branch. The duals to all condition codes are included within single-instruction multiple data image processor 300 which makes the "if any" function simple to code.

As previously described, each execution datapath 358a-n within single-instruction multiple-data image processor 300 is equipped with large local memory 362a-n. Each execution unit 360a-n of each respective execution datapath 358a-n directly accesses its own local memory 362a-n by way of a respective dual port 361a-n or program port 361a-n. Each dual port 361a-n or local memory port 361a-n of image processor 300 is provided with an both A port and a B port. Different signals may be transmitted between each execution unit 360a-n and its local memory 362a-n simultaneously by way of the A and B ports under the control of the program being executed within execution units 360a-n. It will be understood that this transfer by way of local memory ports 361a-n is distinguished from

transfers by way of transfer ports 363a-n under the control of block transfer controller 368.

It will be understood that this type of access to local memories 362a-n by execution units 360a-n involves writing of pointers only. Thus these operations are not actually random accessing of local memories 362a-n. This ability of block transfer controller 368 permits split phase transactions as shown in processing/memory time line 500. These split phase transactions are completely independent of instruction sequencer 352. Thus, block transfer controller 368 operates as a separate instruction engine not directly controlled by instruction sequence controller 352. This allows efficient access to memory for the instruction cache. It permits the cache to be filled quickly even if another block is getting instructions from external memory 364, 366. Therefore, block transfer controller 368 minimizes idling or wait stating of execution datapaths 358a-n while waiting for instructions.

It will be understood by those skilled in the art that conventional imaging processing systems usually provide processor consistency wherein instructions are executed in the order that they are requested from memory. It will also be understood that

single-instruction multiple-data image processor 300 of the present invention is provided with weak processor consistency because block transfer controller 368, functioning as a separate instruction engine, can cause certain memory read requests to pass other memory requests.

As also previously described, access to global memory 366 is shared by all execution datapaths 358a-n within single-instruction multiple-data image processor 300. However, global memory 366 is not directly accessed by execution units 360a-n. Rather, execution units 360a-n access global memory 366 by way of memory interface 370 and global port 376 under the control of block transfer controller 368 and control line 378. Thus a selected port 363a-n may be coupled to global memory 366 by way of port 376. Furthermore, external global memory 366 may be global only in a conceptual sense within image processor 300. All that is required is that any execution datapath 358a-n can read or write any location in external global memory 366 by way of its external global memory port 363 and that external global memory 366 be external to execution datapaths 360a-n.

Within single-instruction multiple-data architecture 300 there is provided an improved method

to more efficiently read blocks of data from system memory 364 and global memory 366 into local memories 362a-n by way of global memory port 363a-n, operate on the data within local memories 362a-n by way of lines 340, 332, 344, and write the results back to global memory 366, again by way of global memory port 363. In order to accomplish these more efficient block read and block write operations, single-instruction multiple-data image processor 300 is provided with block transfer instructions and block transfer architecture. The instructions and architecture are adapted to optimize the movement of data between local execution units 360a-n, local memories 362a-n and global memory 366. These input/output operations within single-instruction multiple-data image processor 300 are handled by autonomous synchronous block transfer controller 368.

Block transfer controller 368 within single-instruction multiple-data image processor 300 allows the transfer of two-dimensional arrays which are conformally displaced. This allows a subblock of a larger image to be copied in a single block operation for example. In general, using source and destination bit maps, conformally displaced blocks may be

transferred even though they do not have the same aspect ratio or alignment in physical memory.

The specification for a block transfer operation initiated by a program within image processor 300 is actually a set of lists of individual block transfers. Each enabled execution datapath 358a-n builds a list of block transfer commands in its local memory 362a-n. A single block transfer initiate instruction eventually leads to the processing of all block transfer commands from the lists of every enabled execution datapath 358a-n. In addition, up to two sets of transfer lists can be pending at any time.

Referring now to Fig. 5, there is shown processing/memory time line 500. As illustrated by processing/memory time line 500, all block transfer operations within single-instruction multiple-data architecture 300 are split-phase transactions that occur concurrently with program execution. That is, a first program may perform a transfer during memory read block time segment 504 and perform processing during processing time segment 502. Overlapping with processing time segment 502, a second program may perform a transfer during memory read block time segment 506.

Referring now to Fig. 6, there is shown a schematic representation of the parameters that specify a memory-to-memory block transfer between global memory 366 and local memories 362a-n. It will be understood by those skilled in the art that while the schematic representation herein sets forth a transfer between global memory 366 and local memories 362a-n, the following discussion applies equally to transfers between system memory 364 and local memories 362a-n. This transfer takes place under the control of block transfer controller 368 within single-instruction multiple-data architecture image processor 300 of the present invention. It is possible to specify, for a single transaction, two-dimensional blocks 602a,b that are part of larger two-dimensional frame arrays 600a,b. Two dimensional frame arrays 600a,b may, respectively, reside within global memory 366 and within a selected local memory 362a-n.

Source frame 600a and source block 602a need not have the same aspect ratio as destination frame 600b and destination block 602b during memory-to-memory block transfers within image processor 300. However, the total number of elements and the width of source block 602a must be equal to the total number of elements and the width of destination block 602b.

In the block transfer request of image processor architecture 300 a set of block transfers is described by encoding the block transfer parameters in local memory 362a-n of each enabled execution datapath 358a-n. The block transfer is started by applying an initiate instruction to block transfer controller 368 by execution datapaths 360a-n by way of transfer control bus 344. It will be understood by those skilled in the art that this applying of the initiate instruction to block transfer controller 368 by way of transfer control bus 344 comprises posting of the transfer.

Referring now to Fig. 7, there is shown block transfer linked request list 700 containing a plurality of block transfer command templates 704, 706, 708 for initiating and specifying block transfers within single-instruction multiple-data image processor 300 of the present invention. Short block transfer command template 708 is an abbreviated form of long block transfer command templates 704, 706. Block transfer controller 368 uses default values for the pitch and width of specified memory blocks in order to permit the use of short command template 708 which specifies less transfer information than long command template 706. After posting a memory transfer, block transfer

controller 368 fetches a command template such as long command template 704 or short command template 708 in order to perform the requested transfer.

If fetched command template 704 is part of a linked list containing further long command template 706 linked to fetched template 704, further template 706 is read after completion of the transfer performed under the direction of fetched template 704. Block transfer controller 368 then performs a further transfer according to further command template 706. Likewise, after the transfer specified by further command template 706 is completed, short command template 708 is fetched because short command template 708 is linked to command template 706.

Execution datapaths 358a-n maintain a status flag for each initiated block transfer operation in order to perform the completion check required for transfers under the control of block transfer controller 368. Each execution datapath 358a-n of image processor 300 may then check for the transfer to be completed by examining the associated status flag maintained by execution datapath 358.

Linked request list 700 within local memory 362a-n of single-instruction multiple data architecture

300 is thus a linked list of block transfer command templates 704, 706, 708. Each command template 704, 706, 708 in linked request list 700 specifies the parameters for a programmed block transfer of data between local memory 362a-n and system memory 364 or global memory 366 of image processor 300. These parameters may be specified explicitly, implicitly or may be determined by default.

For example, internal address field 720 of linked request list 700 may specify a starting address 604 of two-dimensional block 602b within two-dimensional frame array 600b. Two-dimensional frame array 600b resides in the same local memory 362a-n as linked request list 700. External address field 722 of long command templates 704, 706, 708 specifies a destination starting external address 606. External address 606 is the starting address of two-dimensional block 602a within two-dimensional frame array 600a.

Long command templates 704, 706 also contain pitch information for permitting block transfer controller 368 to perform memory-to-memory transfers within single-instruction multiple-data image processor 300. External pitch field 724 of long command templates 704, 706 specifies the external pitch of two-dimensional frame array 600a. Internal pitch field 728

of long command templates 704, 706 specifies the internal pitch of two-dimensional frame array 600b. Short command template 708 is not provided with external pitch field 724 or internal pitch field 728. The width of both two-dimensional blocks 602a,b is stored in pitch field 728.

Link address field 730 of long command template 720 points to long command template 706. The link address field 730 of long command template 706 in turn points to short command template 708. Note that, in addition to external pitch field 704 and internal pitch field 728, width field 726 of long command templates 704, 706 is not present in short command template 708.

Thus, linked request list 700 within image processor 300 is a specification for a series of individual block transfers by block transfer controller 368. Linked request list 700 of each datapath 358a-n is constructed by its respective execution unit 360a-n in its respective local memory 362a-n. The links in linked request list 700 point to the next template 704, 706, 708 in list 700. Templates 704, 706, and 708 are all resident within local memory 362a-n of the same execution datapath 358a-n. A transfer list 700 may be terminated by setting the address in link field 730 of

the last valid command template to some suitable end-of-list indication.

Each enabled execution datapath 358a-n of architecture 300 supplies its linked request list 700 to block transfer controller 368 for a transaction. There may be up to two block transfer request lists 700 simultaneously initiated by each execution datapath 358a-n. In general the number of linked request lists 700 in a transaction controlled by block transfer controller 368 may be up to two times the number of enabled execution datapaths 358a-n.

Memory address 702 of first command template 704 in linked transfer list 700 within local memory 358a-n is always located at the root pointer register of an associated execution unit 360a-n. A microinstruction that posts a block transfer identifies either first linked list 702 or second linked list 710 as corresponding to the root pointer for the transfer. Each enabled execution datapath 358a-n of image processor 300 has its own valid transfer list 700 in place when a transfer is posted.

All of the block transfers for an initiate instruction under the control of block transfer controller 368 within image processor 300 must be

completed before any transfer is processed for any later initiate instruction. This first-in, first-out ordering is essential to maintain the sequential semantics of the single instruction stream of single-instruction multiple-data image processor 300 of the present invention. Thus, if one initiate posts a write to system memory 364 or global memory 366 and the next initiate posts a read of system memory 364 or global memory 366, it is guaranteed that, regardless of the order in which each of the individual datapaths 358a-n are seized, all writes are finished before any read begins. This rule requires only that sets of transfers be initiated and completed in order. This rule does not preclude buffering multiple requests.

As previously described, each local memory 362a-n within single-instruction multiple-data image processor 300 is provided with dual port architecture. One port of the dual port architectural is global memory port 363a-n or external memory port 363a-n. Global memory port 363a-n of each local memory 362a-n is dedicated to data transfers between local memory 362a-n and system memory 364 or global memory 366. Global memory ports 363a-n are formed by transfer control lines 340, 342, and 344.

The other port of the dual port architecture of each local memory 362a-n of image processor 300 is local memory port 361a-n. Local memory port 361a-n is dedicated to transfers between local memories 362a-n and execution units 360a-n. Local memory port 361a-n of the dual port architecture of system 300 is formed of two separate ports, an A port and a B port. Transmission of data by way of the program port of local memory ports 361a-n is under the control of the instructions issued by sequence controller 352 having instruction memory 380. As a result, single-instruction multiple-data architecture 300 can support access to global memory 366, including block transfers, while simultaneously continuing execution within execution units 360a-n.

For example, memory read block time segment 506 of processing/memory time line 500 is simultaneous with processing time segment 502 for its entire duration. This is accomplished by converting the read and write operations of single-instruction multiple-data image processor 300 into split phase transactions comprising an initiate operation and a transfer complete synchronization. This capability is essential for sustaining the high external global memory

bandwidth required by video signal processing applications.

Thus an initiate instruction only initiates the block transfer operation while sequence controller 352 continues to issue instructions to execution units 360a-n while the transfer takes place. It will be understood by those skilled in the art that a program executing within execution units 360a-n of image processor 300 must resynchronize with global memory 366 or external memory 366 after completion of the transfer by block transfer controller 368.

In order to simplify the following example, the instruction sequence of Table IX is written for image processor 300 having a single execution datapath 358a.


```

/* wait for the first block to be read into temp_block1 */
start_read_block(original_image[0],temp_block1);
io_wait();

/* initiate read of next block into temp_block2 */
start_read_block(original_image[1],temp_block2);

for (k = 0; k < NUMBERSOFBLOCKS; k = k + 2);
/* transform two blocks per iteration */
DCT_block(temp_block1);
io_wait();
start_write_block(xform_image[k],temp_block1);
start_read_block(original_image[k+2],temp_block1);

DCT_block(temp_block2);
io_wait();
start_write_block(xform_image[k+1],temp_block2);
start_read_block(original_image[k+3],temp_block2);
);

```

Table IX

The instruction sequence of Table IX uses a straight-forward double buffering technique. In this double buffering technique, while temp_block1 within local memory 362a is transformed, temp_block2 within local memory 362a is loaded with the next block from global memory 366. After the transform of the first block is completed, a write of the transformed data of block1 from local memory 362a to global memory 366 is initiated. This is followed by a read of the next data from global memory 366 for block1. Execution of the instruction sequence of Table IX then advances to transforming temp_block2 simultaneously with writing the results of temp_block1 back to global memory 366.

The new read of temp_block2 from global memory 366 proceeds in parallel with this write of temp_block1.

The io_wait() of Table IX causes sequence controller 352 to wait until all input/output transfers have been completed within execution datapaths 358a-n before proceeding. This guarantees that a previously initiated read or write by execution unit 360a-n is no longer using the source block and that the contents of the destination block are valid. In general, the io_wait() is implemented in image processor 300 with a conditional branch instruction which tests the block transfer done flag local to each execution datapath 358a-n.

As previously described, processing/memory time line 500 is a representation of the processing and memory accesses of the multibuffered input/output of single-instruction multiple-data image processor 300. The latency of input/output transfers are hidden by the overlapped computation on the blocks. For example, consider the blockwise discrete cosine transform of an image within the architecture of single-instruction multiple-data image processor 300. The basic discrete cosine transform algorithm has considerable parallelism. The image to be transformed is tiled by a set of uniformly sized blocks and the discrete cosine

transform is independently computed over each of the blocks. Because there are no data dependencies among the individual block transforms, all sets of block transforms may be computed in parallel by parallel execution datapaths 358a-n.

Referring now to Fig. 8, there is shown single-instruction multiple-data image processor 800 having four execution datapaths 806a-d. Single-instruction multiple data image processor 800 is an alternate embodiment of image processor 300. Sequence controller 802 of image processor 800 applies instructions from instruction memory 801 to execution units 814a-d of execution datapaths 806a-d by way of instruction broadcast line 804. When performing block transfer and block transform operations, single-instruction multiple-data architecture 800 has the simplification of a single instruction stream as previously described with respect to single-instruction multiple-data image processor 300.

The instruction sequence of Table X may be provided within instruction memory 801 to program image processor 800 to perform a blockwise discrete cosine transform of original image memory block 818 stored in global memory 816 to provide a transform image stored in transformed memory image block 820. As also

previously described, the instruction sequence of Table X is executed simultaneously by all execution datapaths 806a-d of four execution datapath image processor 800.

The instruction sequence of Table X may perform the discrete cosine transform required within the instruction sequence of Table IX. The discrete cosine transform instruction sequence of Table X is sequential except that the outer loop *j* is executed only one-fourth as many times as the inner loop *k*. This is due to the four-fold parallelism of execution datapaths 806a-d. The loop induction variable, *j*, of the outer loop takes the values 0,4,8....

DCT_Image:

```
global int original_image[ ], xform_image [ ];
local int j, k, temp_block[BLOCKSIZE];
for (j = 0; j <NUMBEROFBLOCKS;j = j + 4){
    k = j + THIS_DP_NUMBER;
    read_block (original_image[k], temp_block);
    DCT_block(temp_block);
    write_block(xform_image[k], temp_block);
};
```

Table X

The value of the constant THIS_DP_NUMBER, which is used to determine the loop induction variable *k*, depends upon which execution datapath 806a-d is performing the operation. This permits each execution datapath 806a-d of image processor 800 to select a

different block number, k , to process by adding the execution datapath-dependent value `THIS_DP_NUMBER` to the value of j . Thus the constant `THIS_DP_NUMBER`, which is unique to each execution datapath 806a-d, is equal to zero for execution datapath 806a, one for execution datapath 806b, etc. Execution datapath 806a therefore processes block numbers $k = 0, 4, 8, \dots$ of original memory image block 818 from global memory 816, while execution datapath 806b processes blocks $k = 1, 5, 9, \dots$ etc.

All execution datapaths 806a-d copy their assigned blocks k into their respective `temp_block` 812a-d, a temporary array in local memory 810a-d of each execution datapath 806a-d. The assigned blocks of original image memory block 818 copied from global memory 816 into respective local memories 810a-d are then transformed by the respective execution units 814a-d of execution datapaths 806a-d. The resulting transformed blocks within each execution datapath 806a-d are then copied back to transformed memory image block 820 within global memory 816. While original image block 818 and transformed image block 820 are in shared global memory 816, the four sets of values of the loop induction variables i and k reside in respective memory locations 822a-d. Memory locations

822a-d for storing the local values of i and k are provided in each local memory 810a-d of each execution datapath 806a-d, along with the local values of temp_block 812a-d.

In addition to the simplicity of a single instruction stream like that of conventional sequential processors (not shown), single-instruction multiple-data image processor 800 provides significant efficiency advantages for those algorithms that may be optimized for single-instruction multiple-data systems. In addition to the economics of sharing single sequence controller 802 and associated instruction memory 801, an important advantage is the synchronization of multiple execution datapaths 806a-d. For example, during the blockwise discrete cosine transform of the instruction sequence of Table X, it is easy to determine when all execution datapaths 806a-d have completed a current set of block transforms. The block transforms are completed when the last instruction in the block transform code has executed.

A very important feature of single-instruction multiple-data image processor 300 is block transfer controller 368. Block transfer controller 368 controls the flow of data within image processor 300. It prioritizes block transfers of data from global

memory 366 to local memories 362a-n when requested by execution units 360a-n. Block transfer controller 368 also generates addresses and controls for transfers from local memories 362a-n to global memory 366 or system memory 364.

The functions performed by block transfer controller 368 during block data transfers of data between global memory 366 and local memories 362a-n include arbitration of transfer requests of competing global memory 366 and execution units 360a-n, address generation and control for two-dimensional block transfers between global memory 366 and local memories 362a-n, control for scalar, first-in first-out, and statistical decoder transfers between local memories 362a-n and global memory 366 and address generation and control for block instruction load following cache miss.

A number of different types of transfers requiring input/output access arbitration by block transfer controller 368 may take place between global memory 366 and local memories 362a-n. These include fetching instructions from global memory 366. Image processor system 300 initializes the process by downloading instructions from system memory 364 to global memory 366. On power up of image processor 300

the instructions are loaded from system memory 364 or global memory 366 into the controller.

The different types of transfer may be prioritized by block transfer controller 368 as follows, proceeding from highest priority to lowest priority: (1) instruction, (2) scalar, first-in, first-out and statistical decoder, and (3) block transfer. Thus block transfer controller 368 not only prioritizes and arbitrates within image processor 300 based upon whether a request is an instruction type of request or a data type of request. Block transfer controller 368 also prioritizes and arbitrates based upon the subtypes of data.

In order to handle higher priority transfers while suspending lower priority transfers, the transfer having higher priority is allowed to gain access to transfer data bus 340 of global memory ports 363a-n at the next available memory cycle boundary. After the completion of the higher priority transfer, the suspended lower priority transfer is resumed. Thus it will be understood that single-instruction multiple-data image processor 300 of the present invention is provided with weak processor consistency. In the weak processor consistency of image processor 300 input/output operations are not necessarily executed in

the order in which they are applied to execution units 360a-n.

Block transfer controller 368 supports multiple levels of suspension. For example, a block transfer within image processor 300 may be interrupted by a scalar access, which may be interrupted by an instruction load. In this case, block transfer controller 368 suspends the block transfer at the next available memory cycle boundary and starts performing the scalar access. When the cache miss occurs, the scalar access is interrupted at the next memory cycle boundary and the block of instructions is fetched. After the instruction fetch is complete, block transfer controller 368 resumes the scalar access starting with the last execution unit 360a-n serviced. Upon completion of servicing of all the scalar accesses posted, the block transfer is resumed.

When execution units 360a-n require data, they assert a request flag on line 340 to block transfer controller 368. Block transfer controller 368 then arbitrates the requests from the competing execution units 360a-n, and grants an execution unit 360a-n access to the bus. The execution unit 360a-n which was granted access by block transfer controller 368 retains control of the bus until completion of the

transaction unless a request of higher priority is posted by a different execution unit 360a-n.

Block transfer controller 368 services all requests posted at the same time by an execution unit 360a-n before starting the request of another execution unit 360a-n. In order to ensure that no execution unit 360a-n is locked out by posting of an intervening transfer, block transfer controller 368 uses a serial polling scheme to service requests from the competing execution units 360a-n. In this serial polling scheme the last execution unit 360a-n serviced by block transfer controller 368 is not eligible for servicing by controller 368 again until outstanding requests of all other execution unit 360a-n have been serviced.

The two-dimensional block transfer method within single-instruction multiple-data image processor 300 is a direct memory access mechanism for moving one or several rectangular images of arbitrary size between global memory 366 and local memories 362a-n. The operation of the two-dimension block transfer by block transfer controller 368 is autonomous with respect to program execution by execution units 360a-n of single-instruction multiple-data image processor 300. Synchronization between the program being executed by execution units 360a-n and the block transfer operation

controlled by block transfer controller 368 is accomplished by two interactions: the transfer request and the completion check.

As previously described, single-instruction multiple-data image processor 300 is provided with a linked list structure 700 to specify block transfer operations to block transfer controller 368. Each such block transfer linked request list structure 700 for specifying transfers to controller 368 within image processor 300 is provided with a root pointer and one or more block transform command templates 704, 706, 708. Two linked request lists 700 are supported for each execution unit 360a-n, each request list 700 of an execution unit 360a-n having its own set of one or more command templates 704, 706, 708. The programmer of image processor 300 constructs and stores linked request lists 700 in local memories 362a-n prior to initiating a transfer under the control of block transfer controller 368. When a block transfer is initiated by an active execution unit 360a-n, only active execution units 360a-n within image processor 368 are serviced.

In order to provide linked request list 700 within local memories 362a-n, a user of single-instruction multiple-data image processor 300 writes to root pointer registers (not shown) within execution datapaths 358a-n. Additionally, the user writes command templates 704, 706, 708 into local memories 362a-n to set up the required block transfer by block transfer controller 368. The root pointer register holds the starting address of first command template 704 in each execution unit 360a-n.

Command template 704, which contains the description of the first two-dimensional transfer, varies from one execution datapath 362a-n to another within image processor 300. Because multiple command templates 704, 706, 708 may be connected in linked request list 700, as previously described, and two linked lists 700 may be supported, the transfer of multiple independent blocks of data between global memory 366 and local memories 362a-n within datapaths 358a-n of image processor 300 under the control of block transfer controller 368 may thus be initiated by a single instruction.

As previously described, there are two types of command templates, long command templates 704, 706 and short command template 708. The information within

short command template 708 is a subset of the information within long command template 704. Within block transfer controller 368 there are three default registers: the default external pitch register, the default internal pitch register and the default width register. When using long command template 704, or long command template 706, block transfer controller 368 uses the pitch and width value specified fields 724, 726 of long command templates 704, 706. When using short command template 708, block transfer controller 368 uses the pitch and width values specified in default registers. The default register contents may be modified by the user of single-instruction multiple-data image processor 300. When the user performs this modification of the default registers, the default register contents may be updated to contain new external pitch parameters, internal pitch parameters and width parameters as specified by the user.

Prior to initiating a block transfer by block transfer controller 368, the user of single-instruction multiple-data image processor 300 must check a block transfer done bit to determine that the previous block transfer is completed. Initiation of a block transfer without checking the previous transfer block done bit

may cause indeterminate results as previously described. After the previous block transfer is completed and the done bit is set, the user may set a block transfer request bit in block transfer controller 368. This is done by a block transfer initiation instruction. The done bit is cleared as soon as the block transfer is initiated. Several active execution units 360a-n may request block transfers and block transfer controller 368 may arbitrate among active execution units 360a-n to start the block transfers.

After a block transfer is initiated, block transfer controller 368 executes the block transfer sequences starting with execution unit 360a. Block transfer controller 368 fetches the root pointer address of first command template 704 from execution unit 360a if execution unit 360a is active. Block transfer controller 368 uses this root pointer address to load first command template 704 of linked request list 700 of execution unit 360a. After command template 704 is loaded, block transfer controller 368 transfers a memory block between local memory 362a and global memory 366 according to the specifications within command template 704.

When the first block transfer is completed, according to command template 704, block transfer

controller 368 uses link address field 730 of command template 704 to fetch next command template 706 of linked request list 700 of execution unit 360a. Execution of the next block transfer according to command template 706 is then performed. Once block transfer controller 368 loads a link address field 728 having a null, the next active execution unit 360b-n is serviced by block transfer controller 368. The root pointer register of next active execution unit 360b-n is used to fetch the first link address.

When block transfer controller 368 completes linked request lists 700 of all active execution units 360a-n within single-instruction multiple-data image processor 300, a block transfer done condition bit is set. When one set of linked request lists 700 within local memories 362a-n of active datapaths 360a-n is completed in this manner, block transfer controller 368 is prepared to service another set of linked request lists 700 when initiated by execution units 360a-n.

Two block transfer done bits are therefore required in the condition code register (not shown) of each execution datapath 360a-n. One block transfer done bit is provide for each of the two possible linked request lists 700 which may be stored in its local memory 362a-n. These two block transfer done bits are

determined by block transfer controller 368. After each transfer by block transfer controller 368 according to linked request lists 700 is finished, the corresponding block transfer done bit in the condition code register of the serviced datapath 360a-n is set by controller 368. The block transfer done bits are cleared by block transfer controller 368 when the following block transfer is initiated by execution units 360a-n. The block transfer done bits are microcode condition codes. The microcode can use a "branch if condition" instruction to generate an IO_WAIT() function when required for a block transfer between global memory 366 and local memories 362a-n within image processor 300.

Block transfer operations within single-instruction multiple-data image processor 300 have lower access priority than instruction fetches and scalar transfers as previously described. Block transfer controller 368 therefore suspends a lower priority block transfer operation requested by an execution unit 360a-n when a transfer operation of a higher priority is requested by a different execution unit 360a-n. Block transfer controller 368 and memory interface 370 first finish any outstanding read or write request for block transfers and save the

parameters of the command template 704, 706, 708 specifying the block transfer in progress.

The parameters saved by block transfer controller 368 when a higher priority transfer is initiated include the next address in external global memory 366 and the next address in local memory 362a-n of the datapath 358a-n being serviced. These parameters are saved in registers (not shown) within block transfer controller 368. Block transfer controller 368 then services the higher priority access. When the higher priority access is complete, block transfer controller 368 retrieves the saved address parameters from its registers and resumes or reinstates the suspended block transfer.

As previously described, only the instruction fetch and the scalar transfer have higher priority than the block transfer within single-instruction multiple-data architecture 300 as controlled by block transfer controller 368. A block transfer in an interrupt routine does not have higher priority than an initiated block transfer. Thus initiated block transfer is serviced first by controller 368. Block transfer controller 368 thus provides interruptible memory transfers between local memories 362a-n and system memory 364 or global memory 366. If a block transfer

is initiated within single instruction multiple-data image processor 300 while a cache miss, scalar transfer, or first-in first-out access is being serviced by controller 368, it must wait until the higher priority request is finished.

Referring now to Fig. 9, there is shown scalar access transfer system 1000 of single-instruction multiple-data image processor 300. Within scalar access system 1000, there are two sets of scalar transfer controls for block transfer control 368. Within each execution unit 360a-n of image processor 300, there are two scalar address registers 1001, 1002 and two scalar data registers 1004, 1006. Scalar data registers 1004, 1006 may be thirty two bits wide. Scalar address register 1001 and scalar data register 1004 within execution unit 360a-n are associated with the first set of transfer controls. Scalar address register 1002 and scalar data register 1006 are associated with the second set. Address registers 1001, 1002 hold the address in external memory 364, 366. Address registers 1001, 1002 may be loaded by either of two buses, the A-Bus or the B-Bus. Each scalar data register 1004, 1006 can be read or written by either the A-Bus or the B-Bus.

For an output scalar transfer, block transfer controller 368 fetches the contents of scalar address register 1001 or scalar address register 1002 as well as data register 1004 or data register 1006. This fetch is performed by way of transfer data bus 340. For an input scalar transfer, block transfer controller 368 fetches scalar address register 1001, 1002, loads to interface 370, and fetches the return data load to data register 1004, 1006 through transfer data bus 340. The scalar access may support byte, half word, and word alignment. Byte and half word data fetched and loaded in scalar data registers 1004, 1006 are right justified.

For a scalar write, scalar data register 1004 or scalar data register 1006 must be written first. The scalar output address register may then be written. Writing to the scalar output register initiates a scalar output transfer request to block transfer controller 368. For the scalar input transfer, the user of image processor 300 writes scalar input address register 1001 or scalar input address register 1002. The scalar done bit is cleared when the scalar transfer is posted. After the scalar transfer is initiated, block transfer controller 368 samples all active execution units 360a-n and processes the scalar

transfers starting with execution unit 360a. An execution unit 360a-n is determined to be active when its execution mask flag is reset. The scalar done bit of an active execution unit 360a-n must be checked by block transfer controller 368 before initiating the next scalar transfer. Writing to scalar address register 1001, 1002 or scalar data register 1004, 1006 when the scalar done bit is not set causes indeterminate results.

Block transfer controller 368 processes scalar transfers from active execution units 360a-n within image processor 300, beginning with execution unit 360a. For outputs, block transfer controller 368 fetches the scalar address and scalar data from the first active execution unit 360a-n and executes a write to global memory 366 through memory interface 370. For inputs, block transfer controller 368 only fetches the scalar address, executes a read of global memory 366 and loads the data which it read from global memory 366 into the scalar data register.

After processing execution unit 360a, if execution unit 360a is active, block transfer controller 368 proceeds to the remaining active execution units 360b-n until all active execution units 360a-n are serviced. After the first set of scalar

transfers is completed, the first scalar done bit is set by block transfer controller 368. When finished with the first set of scalar transfers, block transfer controller 368 processes the second set in the same sequence. When the second set of scalar transfers is completed, block transfer controller 368 sets the second scalar done bit. Each execution datapath 358a-n of single-instruction multiple-data image processor 300 may have up to two sets of scalar transfers pending at same time within execution datapaths 358a-n. Image processor 300 cannot initiate a third scalar transfer before the first transfer is completed.

The two scalar done bits are located in the condition code register of each execution datapath 358a-n. These two scalar done bits are determined by block transfer controller 368, one for each set of scalar transfers, as previously described. The scalar done bits are cleared when execution units 360a-n again initiate scalar transfers in accordance with instructions applied to execution units 360a-n by sequence controller 352 by way of instruction broadcast line 356. To generate the correct results for scalar transfers, it is therefore necessary to check the scalar done bits before initiating the scalar transfer.

As previously described, the scalar transfer has lower access priority than the instruction fetch within single-instruction multiple-data image processor 300 of the present invention. Block transfer controller 368 of image processor 300 therefore suspends any scalar transfer operation in progress when an instruction fetch is requested. To service the instruction fetch, block transfer controller 368 and memory interface 370 finish the outstanding read or write access of the scalar transfer being suspended. The higher priority instruction fetch is serviced by block transfer controller 368. The suspended scalar transfer is then completed.

Block transfer controller 368 supports scalar internal broadcast transfer. When the microcode writes to the internal broadcast transfer registers the broadcast scalar transfer is initiated. The data in the scalar data register is broadcasted to other scalar data registers. The execution mask bit of all but one of the n individual execution units 360a-n is set in order to initiate a scalar internal broadcast transfer within single-instruction multiple-data image processor 300. The data register of the single unmasked execution unit 360a-n is the source of the request for the scalar internal broadcast transfer. The data

registers of the remaining masked execution units 360a-n are destinations for the scalar internal broadcast transfer. The scalar broadcast only supports word data transfer. The scalar address registers are ignored in the scalar broadcast transfer.

Block transfer controller 368 also supports instruction fetch for single-instruction multiple-data image processor 300. When the instruction cache of single-instruction multiple-data image processor 300 misses within image processor 300, sequence controller 352 is halted. Following this, block transfer controller 368 fetches the external instruction block address from sequence controller 352. Block transfer controller 368 uses the block address to read an instruction block from memory and load it into instruction random access memory 380 located in controller 352. After the instructions are loaded, sequence controller 352 resumes operation and applies instructions to execution datapaths 358a-n.

Referring now to Fig. 10, there is shown a more detailed block diagram representation of statistical decoder input channel 372 of single-instruction multiple-data image processor 300 of the present invention. Statistical decoder input channel 372 of image processor 300 is a specialized input

channel having input channel processor 1116 that reads a variable-length bit sequence from global memory 366 and converts it into a fixed-length bit sequence that is read by execution datapath 358a only.

Because input channel 372 or statistical decoder 372 is a specialized hardware channel having its own input channel processor 1116 it may function as a semi-autonomous unit capable of performing part of the decoding process. This removes some of the burden of decoding input data from execution units 360a-n and eliminates the need for some instruction coding. Input channel processor 1116, execution units 360a-n, and transmission output line 1112 are all formed on a single integrated circuit chip.

Statistical decoder channel 372 functions by having input channel processor 1116 prefetch and decode data. This function is performed when the program executing on executing units 360a-n provides the address of the data to channel processor 1116 and instructs it to begin. It will be understood that the decoding performed by input channel processor 1116 of statistical decoder input channel 372 may take place simultaneously with execution of program instructions by execution units 360a-n. The program executing on execution units 360a-n of image processor 300 later

reads and processes the data received from statistical decoder 372 by way of transmission line 1116.

During image compression by an image processor such as single-instruction multiple-data image processor 300, as well as during other applications such as text compression, certain values within the data being compressed occur more frequently than others. One way to compress data of this nature is to use fewer bits to encode more frequently occurring values and more bits to encode less frequently occurring values. This type of encoding results in a variable-length sequence in which the length of a code may range, for example, from one bit to sixteen bits. It will be understood by those skilled in the art that a code is a group of bits used to encode a single value.

Statistical decoder input channel input channel 372 includes get_next_8-bits logic block 1104 and memory 1108. Memory 1108 may store a decoding tree such as a conventional Huffman coding tree. Decoder 372 may, for example, decode up to eight bits in code length. Input channel processor 1116 of statistical decoder 372 determines the next eight parallel bits of a bitstream applied to decoder 372 by way of input line 1102 and uses them as an address on address line 1106

to access memory 1108. Using the address information in this way, statistical decoder 372 obtains a decoded value from memory 1108. This decoded memory output value, which may have a twenty bit format, appears on memory output line 1110.

In the preferred embodiment, four bits of the format stored in memory 1108 may be used to indicate the size of the bits to be shifted in the next cycle, one bit may provide a flag for indicating when the decoding is completed, and fifteen bits may represent the value returned to the microcode executing within execution datapaths 358a-n. Feedback loop 1114 feeds the four bits of shift size back to get_next_8-bit logic block 1104. In the preferred embodiment the sixteen bit output value appears on decoder output line 1112.

Statistical decoder input channel 372 of single-instruction multiple-data processor 300 may operate in one of two modes, a native mode and a compatible mode. In the native mode, statistical decoder 372 may be used for Huffman decoding. In this mode of statistical decoder 372, decoder 372 decodes data provided by a conventional Huffman coding scheme in a manner well known to those skilled in the art. The Huffman coding scheme is described in Huffman,

D.A., "A Method for the Construction of Minimum Redundancy Codes," Proc. Inst. Electr. Radio Eng. 40, 9, September, 1952, pages 1098-1101.

When in the compatible mode, statistical decoder 372 of image processor 300 may be used for the statistical decoding of data encoded by coding schemes other than a conventional Huffman encoding scheme in accordance with the operation of compatible image processors (not shown). Operation of statistical decoder 372 in the native mode or in the conditional mode may be selected by writing an appropriate value into a control register (not shown) within sequence controller 352. It is possible to use statistical decoder 372 in either of these two modes, whether image processor 300 is running in its own native mode or in a compatible mode provided for processing of data in accordance with a compatible image processor (not shown).

Referring now to Fig. 11, there is shown binary decoding tree 1200. Binary decoding tree 1200 represents a coding scheme such as a conventional Huffman coding scheme. Binary decoding tree 1200 is used by statistical decoder 372 to convert variable-length bit input sequences into fixed length bit sequences within single-instruction multiple-data image

processor 300. Memory 1108 of statistical decoder 372 may store binary decoding tree 1200 for decoding variable length coded data within single-instruction multiple-data image processor 300 of the present invention. Access to values determined in accordance with binary decoding tree 1200 and stored within memory 1108 is obtained by way of address line 1106 as previously described.

Decoding by means of binary decoding tree 1200 begins at root node 1202. When decoding is performed bit-by-bit, statistical decoder 372 tests the next bit from the bitstream of decoder input line 1102 to determine whether the next bit has the value one or the value zero. Statistical decoder 372 takes right branch 1204 of binary decoding tree 1200 if the tested bit is one. Statistical decoder 372 takes left branch 1206, if the tested bit is zero.

There is a `decode_completed` flag associated with each node of binary decoding tree 1200. The `decode_completed` flag indicates whether the decoding of a bit sequence is completed at the node. If the `decode_completed` flag for a node is set, statistical decoder 372 stops decoding, and reports the value stored in the node. Then, the next decoding starts

from root node 1202. If the flag is not set, statistical decoder 372 tests the next bit of the input bitstream and continues from its current node.

In this example, decoding is performed by statistical decoder 372 bit-by-bit while using binary decoding tree 1200. However, statistical decoder 372 of image processor 300 may also decode several bits of the input bitstream at a time.

It will be understood that various changes in the details, materials and arrangements of the parts which have been described and illustrated in order to explain the nature of this invention may be made by those skilled in the art without departing from the principle and scope of the invention as expressed in the following claims.

CLAIMS

1. A data processing system having executing means for executing an instruction sequence, comprising:

means for sequentially determining at least first and second conditionals in accordance with instructions of said instruction sequence;

means for setting respective first and second mask flags in accordance with said first and second determined conditionals;

means for storing said first and second mask flags;

means for sequentially retrieving said first and second mask flags in a predetermined order; and,

said executing means having means for executing selected instructions of said instruction sequence in accordance with said sequentially retrieved mask flags.

2. The data processing system of Claim 1, wherein said executing means comprises means for idling

during selected instruction cycles in accordance with said retrieved mask flags.

3. The data processing system of Claim 1, further comprising means for storing said first and second mask flags sequentially.

4. The data processing system of Claim 3, wherein said sequential storing means comprises a stack.

5. The data processing system of Claim 3, comprising means for storing a further mask flag after retrieving one mask flag of said first and second sequentially stored mask flags and before retrieving the remaining mask flag of said first and second mask flags.

6. The data processing system of Claim 5, wherein said means for sequentially retrieving comprises means for retrieving said further mask flag and said remaining mask flag of said first and second mask flags in a further predetermined order.

7. The data processing system of Claim 1, further comprising instruction sequencing means for applying instructions of said instruction sequence to said executing means, said instruction sequencing means

including means for withholding selected instructions from said executing means during selected instruction cycles.

8. The data processing system of Claim 7, further comprising mask flag determining means for selecting said selected instruction cycles in accordance with said determined mask flags.

9. The data processing system of Claim 1, wherein said second conditional is determined in response to said determining of said first conditional.

10. The data processing system of Claim 1, wherein there is provided external memory and a plurality of said executing means arranged as parallel datapaths for simultaneously executing an identical instruction of said instruction sequence, each executing means having at least one individual mask flag and mask flag setting means, comprising:

means for determining the respective mask flags of each of said executing means; and,

means for executing an instruction in selected executing means and simultaneously idling in other executing means in accordance with said respective mask flags.

11. The data processing system of Claim 10, wherein said means for determining said respective mask flags further comprises:

consensus means for determining that the respective mask flags of all executing means of said plurality of executing means are in the same state; and,

means for withholding instructions of said instruction sequence from all of said executing means in response to said consensus determination.

12. The data processing system of Claim 10, wherein each of said executing means is provided with a respective local memory having a dual port.

13. The data processing system of Claim 12, wherein each local memory is accessed only by its respective execution means.

14. The data processing system of Claim 12, wherein said dual port comprises a local port for coupling said local memory to its respective executing means and a global port for coupling said local memory to said external memory.

15. The data processing system of Claim 14, wherein said local port transmits data simultaneously

with transmission of data by way of said external port.

16. The data processing system of Claim 15, further comprising data transfer control means for controlling said external port and said simultaneous transfer of data.

17. The data processing system of Claim 16, wherein said data transfer control means comprises block transfer control means for controlling transmission of blocks of data.

18. The data processing system of Claim 16, wherein said data transfer control means is external to said execution means.

19. A single-instruction multiple-data architecture having external memory and at least first and second datapaths, comprising:

respective first and second execution means within said first and second datapaths respectively for parallel processing of data within said first and second datapaths;

respective first and second datapath local memories within said first and second datapaths respectively;

respective first and second local memory ports for coupling each of said first and second execution means directly to said external memory and to its respective datapath local memory; and,

said coupling of said first and second local memory local memory ports being adapted to permit each of said first and second execution means direct local memory access only to its respective first and second local memories.

20. The single-instruction multiple-data architecture of Claim 19, wherein said local memory ports comprise a local port for coupling said local memory to its respective executing means and a global port for coupling said local memory to said external memory.

21. The single-instruction multiple-data architecture of Claim 20, wherein said local port transmits data simultaneously with transmission of data by way of said global port.

22. The single-instruction multiple-data architecture of Claim 21, further comprising transfer controller means for controlling said global port and said simultaneous transmission of data.

23. The single-instruction multiple-data architecture of Claim 22, wherein said transfer controller means comprises block transfer controller means for controlling the transfer of blocks of data.

24. The single-instruction multiple-data architecture of Claim 23, wherein said block transfer controller means is external to said first and second executing means.

25. The single-instruction multiple-data architecture of Claim 19, further comprising sequence controller means for simultaneously applying a single instruction to both said first datapath and to said second datapath for simultaneous execution of said applied instruction by said first and second executing means.

26. A video processor semiconductor device having semiconductor device transmission means within said device and at least one datapath, comprising:

first processing means within said datapath for first processing data within said datapath;

input channel means coupled to said datapath for receiving coded channel input data;

said input channel means having second processing means, separate from said first processing means, for second processing and decoding said received coded channel input data simultaneously with said first processing to provide decoded channel output data; and,

said semiconductor device transmission means being adapted to receive said decoded channel output data from said second separate processing means and apply said decoded channel output data to said first processing means within said datapath.

27. The video processor semiconductor device of Claim 26, wherein said second separate processing means comprises statistical decoding means.

28. The video processor semiconductor device of Claim 27, wherein said statistical decoding means comprises means for decoding a variable length code.

29. The video processor semiconductor device of Claim 26, wherein said second separate processing means comprises:

memory means having memory locations for storing channel output data; and,

means for accessing selected memory locations within said memory means in accordance with at least a portion of said coded channel input data.

30. A method for processing data in a data processing system having executing means for executing an instruction sequence, comprising the steps of:

(a) sequentially determining at least first and second conditionals in accordance with instructions of said instruction sequence;

(b) setting respective first and second mask flags in accordance with said first and second determined conditionals;

(c) storing said first and second mask flags;

(d) sequentially retrieving said first and second mask flags in a predetermined order; and,

(e) executing selected instructions of said instruction sequence in accordance with said sequentially retrieved mask flags.

31. The method for processing data of Claim 30, wherein step (e) further comprises the step of

idling said executing means during selected instruction cycles in accordance with said retrieved mask flags.

32. The method for processing data of Claim 30, wherein step (c) further comprises the step of storing said first and second mask flags sequentially.

33. The method for processing data of Claim 32, wherein comprises step (c) further comprises the step of sequentially storing said first and second mask flags in a stack.

34. The method for processing data of Claim 32, comprising the further step of storing a further mask flag after retrieving one mask flag of said first and second sequentially stored mask flags and before retrieving the remaining mask flag of said first and second mask flags.

35. The method for processing data of Claim 34, further comprising the step of retrieving said further mask flag and said remaining mask flag of said first and second mask flags in a further predetermined order.

36. The method for processing data of Claim 30, wherein said system is provided with

executing means and instruction sequencing means for applying instructions of said instruction sequence from said instruction sequencing means to said executing means, comprising the further step of withholding selected instructions from said executing means during selected instruction cycles by said instruction sequencing means.

37. The method for processing data of Claim 36, wherein said system is provided with mask determining means comprising the further step of selecting said selected instruction cycles in accordance with said determined mask flags.

38. The method for processing data of Claim 30, comprising the further step of determining said second conditional in response to said determining of said first conditional.

39. The method for processing data of Claim 31, wherein said data processing system is provided with external memory and a plurality of executing means arranged as parallel datapaths for simultaneously executing an identical instruction of said instruction sequence, each executing means having at least one individual mask flag and mask flag setting means, comprising the further steps of:

(f) determining a conditional within each respective executing means of said plurality of executing means in accordance with a value local to each of said executing means;

(g) setting said individual mask flags by said individual mask flag setting means in accordance with said locally determined conditionals;

(h) determining the individual mask flags of each of said executing means; and,

(i) executing an instruction in selected executing means and simultaneously idling in other executing means in accordance with said individual mask flags.

40. The method for processing data of Claim 39, wherein step (h) comprises the further steps of:

(j) consensus determining that the individual mask flags of all executing units of said plurality of executing means are in the same state; and,

(k) withholding instructions of said instruction sequence from all of said executing means in response to said consensus determination.

41. The method for processing data of Claim 39, wherein each of said executing means is provided with a respective local memory having a dual port.

42. The method for processing data of Claim 41, further comprising the step of accessing each local memory only by its respective execution means.

43. The method for processing data of Claim 41, wherein said dual port comprises a local port for coupling said local memory to its respective executing means and a global port for coupling said local memory to said external memory.

44. The method for processing data of Claim 43, further comprising the step of transmitting data by way of said local port simultaneously with transmission of data by way of said external port.

45. The method for processing data of Claim 44, comprising the further step of providing control means for controlling said external port and said simultaneous transmission of data.

46. The method for processing data of Claim 45, wherein said data transfer control means

comprises block transfer control means for controlling transmission of blocks of data.

47. The method for processing data of Claim 45, wherein said data transfer control means is external to said executing means.

48. A method for processing data in a system having a single-instruction multiple-data architecture having external memory and at least first and second datapaths, comprising:

(a) parallel processing of data by first and second executing means within said first and second datapaths respectively;

(b) providing first and second datapath local memories within said first and second datapaths respectively;

(c) providing first and second local memory dual ports within said first and second datapaths respectively, each dual port of said first and second dual ports having an external memory port and executing means port;

(d) first coupling said external memory ports of said first and second dual ports to said

external memory for providing access between said external memory and each datapath local memory;

(e) second coupling said executing means port of each of said first and second dual ports to said first and second executing means respectively for providing access between each of said executing means and its respective datapath local memory; and,

(f) adapting said first and second local memory ports to permit access to each of said first and second datapath local memories only by its respective executing means.

49. The method for processing data of Claim 48, comprising the further step of transmitting data by way of said external memory port simultaneously with transmission of data by way of said executing means port.

50. The method for processing data of Claim 49, comprising the further step of providing transfer controller means for controlling said executing means port and said simultaneous transfer of data.

51. The method for processing data of Claim 50, wherein said transfer controller means comprises block transfer controller means for controlling the transfer of blocks of data.

52. The method for processing data of Claim 50, wherein said data transfer controller means is external to said first and second executing means.

53. The method for processing data of Claim 48, comprising the further step of transferring system data between said global memory means and each of said first and second datapath local memories for processing of said system data within respective datapaths.

54. The method for processing data of Claim 48, comprising the further step of simultaneously applying a single instruction to both said first datapath and said second datapath for simultaneous execution of said applied instruction by said first and second executing means.

55. A method for processing video data on a single semiconductor device having semiconductor device transmission means within said device and at least one datapath, comprising the steps of:

(a) first processing of data by first processing means within said datapath;

(b) coupling input channel means to said datapath for receiving coded channel input data;

(c) second processing said received coded channel input data simultaneously with said first processing to provide decoded channel output data by second processing means within said input channel means separate from said first processing means;

(d) transmitting said decoded channel output data by way of said semiconductor device transmission means from said second separate processing means to said first processing means within said datapath.

56. The method for processing video data of Claim 55, wherein said second separate processing means comprises statistical decoding means.

57. The method for processing video data of Claim 56, wherein the step of second processing comprises decoding a variable length code.

58. The method for processing video data of Claim 55, wherein step (c) further comprises the steps of:

(e) providing memory means having memory locations for storing said channel output data; and,

(f) accessing selected locations within said memory means in accordance with at least a portion of said coded channel input data.

59. A data processing system having at least one datapath, comprising:

first processing means within said datapath for executing instructions of an instruction sequence and providing memory request signals in accordance with said instructions;

local memory means within said datapath;

first port means coupled to said executing means and to said local memory means for transferring first signals between said executing means and said local memory means;

arbitrating means for receiving said memory request signals and determining a priority memory request signal; and,

second processing means, separate from said first processing means, for controlling said first

port means in accordance with said priority memory request signal.

60. The data processing system of Claim 59, wherein said memory request signals comprise a plurality of differing memory request types.

61. The data processing system of Claim 60, wherein a first memory request type of said plurality of memory request types comprises instruction request signals and a second memory request type of said plurality of memory request types comprises data request signals.

62. The data processing system of Claim 60, wherein said arbitrating means further comprises:

means for determining the memory request types of said received memory request signals; and,

means for determining said priority memory request signal in accordance with said determined memory request types to provide arbitrated priorities among said memory request types.

63. The data processing system of Claim 62, wherein at least one memory request type of said plurality of differing memory request types is provided with a plurality of differing subtypes.

64. The data processing system of Claim 63, wherein said arbitrating means further comprises:

means for determining the memory request subtype of said memory request signals; and,

means for determining said priority memory request signal in accordance with said determined memory request subtype to provide arbitrated priorities among said subtypes.

65. The data processing system of Claim 63, wherein said differing memory request subtypes comprise scalar request signals.

66. The data processing system of Claim 65, wherein said differing memory request subtypes include block transfer request signals.

67. The data processing system of Claim 59, further comprising second port means coupled to said executing means and to said local memory means for transferring second signals between said executing means and said local memory means.

68. The data processing system of Claim 67, wherein said second processing means comprises means for controlling said second port in accordance with said priority memory request signal to permit transfer

of said first signals through said first port means and simultaneous transfer of said second signals through said second port means.

69. The data processing system of Claim 68, wherein said first signals comprise instruction signals and said second signals comprise data signals.

70. The data processing system of Claim 59, including a plurality of said datapaths, further comprising:

each datapath having individual executing means, individual local memory means, and individual first port means coupled between said individual executing means and said individual local memory means; and,

said second processing means having means for controlling said first port means of each datapath in accordance with said priority memory request signal.

71. The data processing system of Claim 70, further comprising:

individual second port means coupled between said individual executing means and said individual local memory of each of said datapaths for

transferring second signals between said executing means and said local memory; and,

said second processing means having means for controlling said second port means of each datapath in accordance with said priority memory request signal.

72. The data processing system of Claim 71, wherein said second processing means comprises means for controlling each of said first and second port means in accordance with said priority memory request signal to permit transfer of said first signals through said first port means simultaneously with transfer of said second signals through said second port means.

73. The data processing system of Claim 59, including interrupt means, further comprising:

means for determining a first priority memory request signal to provide first control of said first port means in accordance with said first priority memory request signal;

means for determining a second priority memory request signal in accordance with said memory request signals while providing said first control of said first port means;

means for suspending said first control to provide second control of said first port means in accordance with said second priority memory request signal; and,

means for reinstating said first control in accordance with said first priority memory request signal after said second control.

74. A method for data processing in a system having local memory means and at least one datapath for executing instructions of an instruction set, comprising the steps of:

(a) executing instructions of said instruction set by first processing means within said datapath;

(b) providing memory request signals in accordance with said executed instructions;

(c) coupling first port means between said executing means and said local memory means for transferring first signals between said executing means and said local memory means;

(d) arbitrating said memory request signals to determine a priority memory request signal; and,

(e) controlling said first port means by second processing means, separate from said first processing means, in accordance with said priority memory request signal.

75. The method of data processing of Claim 74, wherein step (b), comprises providing at least first and second memory request types.

76. The method of data processing of Claim 75, wherein said first memory request type comprises instruction request signals and said second memory request type comprises data request signals.

77. The method of data processing of Claim 75, further comprising the steps of:

(f) determining the memory request types of said memory request signals; and,

(g) determining said priority memory request signal in accordance with said determined memory request types to provide priorities among said memory request types.

78. The method of data processing of Claim 75, wherein at least one of said first and second memory request types is provided with a plurality of differing memory request subtypes.

79. The method of data processing of Claim 77, comprising further steps of:

(h) determining the memory request subtype of said memory request signals within one of said determined memory request types; and,

(i) determining said priority memory request signal in accordance with said determined memory request subtype to provide prioritization of said memory request subtypes.

80. The method of data processing of Claim 78, wherein said differing memory request subtypes include data request scalar request signals.

81. The method of data processing of Claim 80, wherein said differing memory request subtypes include block transfer request signals.

82. The method of data processing of Claim 74, comprising the further step of coupling second port means to said executing means and to said local memory means for transferring second signals between said executing means and said local memory means.

83. The method of data processing of Claim 82, comprising the further steps of:

(j) controlling said second port means by said second processing means in accordance with said priority memory request signal; and,

(k) simultaneously transferring said first signals through said first port means and said second signals through said second port means.

84. The method of data processing of Claim 83, wherein said first signals comprise instruction signals and said second signals comprise data signals.

85. The method of data processing of Claim 74, comprising the further steps of:

(l) providing a plurality of said datapaths, each datapath having individual executing means, individual local memory means, and individual first port means coupled between said individual executing means and said local memory means; and,

(m) controlling each of said individual first port means of each datapath by said second processing means in accordance with said priority memory request signal.

86. The method of data processing of Claim 85, comprising the further steps of:

(n) coupling respective second port means between said individual executing means and said individual local memory means of each of said datapaths for transferring second signals between said individual executing means and said individual local memory means; and,

(o) controlling said second port means of each datapath by said second processing means in accordance with said priority memory request signal.

87. The method of data processing of Claim 86, further comprising the step of controlling said first and second port means to permit simultaneous transfer of said first signals through said first port means and said second signals through said second port means.

88. The method of data processing of Claim 74, wherein steps (d) and (e) comprise the further steps of:

(p) determining a first priority memory request signal to provide first control of said first port means in accordance with said first priority memory request signal;

(q) determining a second priority memory request signal while providing said first control of said first port means;

(r) suspending said first control to provide second control of said first port means in accordance with said second priority memory request signal; and,

(s) reinstating said first control in accordance with said first priority memory request signal after ending of said second control to provide interruptible transfer of said first signals between said executing means and said local memory.




FIG. 1
Prior Art




FIG. 2
Prior Art




FIG. 3




FIG. 4A




FIG. 4B




FIG. 4C




FIG. 4D




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

